

Design Considerations

This chapter describes important design considerations and provides an example of a wireless mesh design.

Each outdoor wireless mesh deployment is unique, and each environment has its own challenges with available locations, obstructions, and available network infrastructure. Design requirements driven by expected users, traffic, and availability needs are also major design criteria. This chapter contains the following sections:

- Wireless Mesh Constraints, on page 1
- ClientLink Technology, on page 5
- Controller Planning, on page 8

Wireless Mesh Constraints

The following are a few system characteristics to consider when you design and build a wireless mesh network. Some of these characteristics apply to the backhaul network design and others to the CAPWAP controller design:

Wireless Backhaul Data Rate

Backhaul is used to create only the wireless connection between the access points. The backhaul interface is 802.11a/n/ac depending upon the access point. The rate selection is important for effective use of the available RF spectrum. The rate can also affect the throughput of client devices, and throughput is an important metric used by industry publications to evaluate vendor devices.

Dynamic Rate Adaptation (DRA) introduces a process to estimate optimal transmission rate for packet transmissions. It is important to select rates correctly. If the rate is too high, packet transmissions fail resulting in communication failure. If the rate is too low, the available channel bandwidth is not used, resulting in inferior products, and the potential for catastrophic network congestion and collapse.

Data rates also affect the RF coverage and network performance. Lower data rates, for example 6 Mbps, can extend farther from the access point than can higher data rates, for example 1300 Mbps. As a result, the data rate affects cell coverage and consequently the number of access points required. Different data rates are achieved by sending a more redundant signal on the wireless link, allowing data to be easily recovered from noise. The number of symbols sent out for a packet at the 1-Mbps data rate is higher than the number of symbols used for the same packet at 11 Mbps. Therefore, sending data at the lower bit rates takes more time than sending the equivalent data at a higher bit rate, resulting in reduced throughput.

A lower bit rate might allow a greater distance between MAPs, but there are likely to be gaps in the WLAN client coverage, and the capacity of the backhaul network is reduced. An increased bit rate for the backhaul

network either requires more MAPs or results in a reduced SNR between MAPs, limiting mesh reliability and interconnection.

Note The data rate can be set on the backhaul on a per AP basis. It is not a global command.

The required minimum LinkSNR for backhaul links per data rate is shown in Table 1: Backhaul Data Rates and Minimum LinkSNR Requirements, on page 2.

Table 1: Backhaul Data Rates and Minimum LinkSNR Requirements

802.11a Data Rate (Mbps)	Minimum Required LinkSNR (dB)
54	31
48	29
36	26
24	22
18	18
12	16
9	15
6	14

- The required minimum LinkSNR value is driven by the data rate and the following formula: *Minimum SNR + fade margin*.

  Table 2: Backhaul Data Rates and Minimum LinkSNR Requirements for 802.11n, on page 2 summarizes the calculation by data rate.

  - Minimum SNR refers to an ideal state of noninterference, nonnoise, and a system packet error rate (PER) of no more than 10 percent.
  - Typical fade margin is approximately 9 to 10 dB.

  Minimum Required LinkSNR Calculations by Data Rate

Table 2: Backhaul Data Rates and Minimum LinkSNR Requirements for 802.11n

802.11n Date Rate (Mbps)	Spatial Stream	Minimum Required LinkSNR (dB)
15	1	9.3
30	1	11.3

802.11n Date Rate (Mbps)	Spatial Stream	Minimum Required LinkSNR (dB)
45	1	13.3
60	1	17.3
90	1	21.3
120	1	24.3
135	1	26.3
157.5	1	27.3
30	2	12.3
60	2	14.3
90	2	16.3
120	2	20.3
180	2	24.3
240	2	27.3
270	2	29.3
300	2	30.3

- If we take into account the effect of MRC for calculating Minimum Required Link SNR. Table 3: Required LinkSNR Calculations for 802.11a/g, on page 3 shows the required LinkSNR for 802.11a/g (2.4 GHz and 5 GHz) for AP1552 and 1522 with 3 Rx antennas (MRC gain).

  LinkSNR = Minimum SNR - MRC + Fade Margin (9 dB)

Table 3: Required LinkSNR Calculations for 802.11a/g

802.11a/g MCS (Mbps)	Modulation	Minimum SNR (dB)	MRC Gain from 3 RXs (dB)	Fade Margin (dB)	Required Link SNR (dB)
6	BPSK 1/2	5	4.7	9	9.3
9	BPSK 3/4	6	4.7	9	10.3
12	QPSK 1/2	7	4.7	9	11.3
18	QPSK 3/4	9	4.7	9	13.3
24	16QAM 1/2	13	4.7	9	17.3
36	16QAM 3/4	17	4.7	9	21.3
48	64QAM 2/3	20	4.7	9	24.3
54	64QAM 3/4	22	4.7	9	26.3

If we consider only 802.11n rates, then Table 4: Requirements for LinkSNR with AP1552 for 2.4 and 5 GHz, on page 4 shows LinkSNR requirements with AP1552 for 2.4 and 5 GHz.

Table 4: Requirements for LinkSNR with AP1552 for 2.4 and 5 GHz

No. of Spatial Streams	11n MCS	Modulation	Minimum SNR (dB)	MRC Gain from 3 RXs (dB)	Fade Margin (dB)	Link SNR (dB)
1	MCS 0	BPSK 1/2	5	4.7	9	9.3
1	MCS 1	QPSK 1/2	7	4.7	9	11.3
1	MCS 2	QPSK 3/4	9	4.7	9	13.3
1	MCS 3	16QAM 1/2	13	4.7	9	17.3
1	MCS 4	16QAM 3/4	17	4.7	9	21.3
1	MCS 5	64QAM 2/3	20	4.7	9	24.3
1	MCS 6	64QAM 3/4	22	4.7	9	26.3
1	MCS 7	64QAM 5/6	23	4.7	9	27.3
2	MCS 8	BPSK 1/2	5	1.7	9	12.3
2	MCS 9	QPSK 1/2	7	1.7	9	14.3
2	MCS 10	QPSK 3/4	9	1.7	9	16.3
2	MCS 11	16QAM 1/2	13	1.7	9	20.3
2	MCS 12	16QAM 3/4	17	1.7	9	24.3
2	MCS 13	64QAM 2/3	20	1.7	9	27.3
2	MCS 14	64QAM 3/4	22	1.7	9	29.3
2	MCS 15	64QAM 5/6	23	1.7	9	30.3

Note With two spatial streams, the MRC gain is halved, that is the MRC gain is reduced by 3 dB. This is because the system has 10 log (3/2 SS) instead of 10 log (3/1 SS). If there were to have been 3 SS with 3 RX, then the MRC gain would have been zero.

- Number of backhaul hops is limited to eight but we recommend three to four hops.

  The number of hops is recommended to be limited to three or four primarily to maintain sufficient backhaul throughput, because each mesh access point uses the same radio for transmission and reception of backhaul traffic, which means that throughput is approximately halved over every hop. For example, the maximum throughput for 24 Mbps is approximately 14 Mbps for the first hop, 9 Mbps for the second hop, and 4 Mbps for the third hop.

- Number of MAPs per RAP.

  There is no current software limitation on how many MAPs per RAP you can configure. However, it is suggested that you limit the number to 20 MAPs per RAP.

- Number of controllers
  - The number of controllers per mobility group is limited to 72.

- Number of mesh access points supported per controller.

ClientLink Technology

Many networks still support a mix of 802.11n and 802.11ac clients. Because 802.11a/g clients (legacy clients) operate at lower data rates, the older clients can reduce the capacity of the entire network. Cisco's ClientLink technology can help solve problems related to adoption of 802.11n in mixed-client networks by ensuring that 802.11a/g clients operate at the best possible rates, especially when they are near cell boundaries.

Advanced signal processing has been added to the Wi-Fi chipset. Multiple transmit antennas are used to focus transmissions in the direction of the 802.11a/g client, increasing the downlink signal-to-noise ratio and the data rate over range, thereby reducing coverage holes and enhancing the overall system performance. This technology learns the optimum way to combine the signal received from a client and then uses this information to send packets in an optimum way back to the client. This technique is also referred to as MIMO (multiple-input multiple-output) beamforming, transmit beamforming, or cophasing, and it is the only enterprise-class and service provider-class solution in the market that does not require expensive antenna arrays.

The 802.11n systems take advantage of multipath by sending multiple radio signals simultaneously. Each of these signals, called a spatial stream, is sent from its own antenna using its own transmitter. Because there is some space between these antennas, each signal follows a slightly different path to the receiver, a situation called spatial diversity. The receiver has multiple antennas as well, each with its own radio that independently decodes the arriving signals, and each signal is combined with signals from the other receiver radios. This results in multiple data streams receiving at the same time. This enables a higher throughput than previous 802.11a/g systems, but requires an 802.11n capable client to decipher the signal. Therefore, both AP and client need to support this capability. Due to the complexity of issues, in the first generation of mainstream 802.11n chipsets, neither the AP nor client chipsets implemented 802.11n transmit beamforming. Therefore, the 802.11n standard transmit beamforming will be available eventually, but not until the next generation of chipsets take hold in the market. We intend to lead in this area going forward.

We realized that for the current generation of 802.11n APs, while the second transmit path was being well utilized for 802.11n clients (to implement spatial diversity), it was not being fully used for 802.11a/g clients. In other words, for 802.11 a/g clients, some of the capabilities of the extra transmit path was lying idle. In addition, we realized that for many networks, the performance of the installed 802.11 a/g client base would be a limiting factor on the network.

To take advantage of this fallow capacity and greatly enhance overall network capacity by bringing 802.11 a/g clients up to a higher performance level, we created an innovation in transmit beamforming technology, called ClientLink.

ClientLink uses advanced signal processing techniques and multiple transmit paths to optimize the signal received by 802.11a/g clients in the downlink direction without requiring feedback. Because no special feedback is required, Cisco ClientLink works with all existing 802.11a/g clients.

Cisco ClientLink technology effectively enables the access point to optimize the SNR exactly at the position where the client is placed. ClientLink provides a gain of almost 4 dB in the downlink direction. Improved SNR yields many benefits, such as a reduced number of retries and higher data rates. For example, a client at the edge of the cell that might previously have been capable of receiving packets at 12 Mbps could now receive them at 36 Mbps. Typical measurements of downlink performance with ClientLink show as much as 65 percent greater throughput for 802.11a/g clients. By allowing the Wi-Fi system to operate at higher data rates and with fewer retries, ClientLink increases the overall capacity of the system, which means an efficient use of spectrum resources.

ClientLink in the 1552 access points is based on ClientLink capability available in AP3500s. Therefore, the access point has the ability to beamform well to nearby clients and to update beamforming information on 802.11ACKs. Therefore, even if there is no dedicated uplink traffic, the ClientLink works well, which is beneficial to both TCP and UDP traffic streams. There are no RSSI watermarks, which the client has to cross to take advantage of this Beamforming with Cisco 802.11n access points.

ClientLink can beamform to 15 clients at a time. Therefore, the host must select the best 15 if the number of legacy clients exceeds 15 per radio. AP1552 has two radios, which means that up to 30 clients can be beamformed in time domain.

Although ClientLink is applied to legacy OFDM portions of packets, which refers to 11a/g rates (not 11b) for both indoor and outdoor 802.11n access points, there is one difference between ClientLink for indoor 11n and ClientLink for outdoor 11n. For indoor 11n access points, SW limits the affected rates to 24, 36, 48, and 54 Mbps. This is done to avoid clients sticking to a far away AP in an indoor environment. SW also does not allow ClientLink to work for those rates for 11n clients because the throughput gain is so minimal. However, there is a demonstrable gain for pure legacy clients. For outdoor 11n access points, we do need more coverage. Thus, three more additional legacy data rates lower than 24 Mbps have been added. ClientLink for outdoors is applicable to legacy data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps.



Note ClientLink is enabled by default.

Configuring ClientLink (CLI)

From the 7.2 release onwards, it is not possible to configure ClientLink (beamforming) using the controller GUI.

Step 1 Disable the 802.11a or 802.11b/g network by entering this command:

config {**802.11a** | **802.11b**} **disable network**

Step 2 Globally enable or disable beamforming on your 802.11a or 802.11g network by entering this command:

config {**802.11a** | **802.11b**} **beamforming global** {**enable** | **disable**}

The default value is disabled.

Note After you enable beamforming on the network, it is automatically enabled for all the radios applicable to that network type.

Step 3 Override the global configuration and enable or disable beamforming for a specific access point by entering this command:

config {**802.11a** | **802.11b**} **beamforming ap** *Cisco_AP* {**enable** | **disable**}

The default value is disabled if beamforming is disabled on the network and enabled if beamforming is enabled on the network.

Step 4 Reenable the network by entering this command:

config {**802.11a** | **802.11b**} **enable network**

Step 5 Save your changes by entering this command:

save config

Step 6 See the beamforming status for your network by entering this command:

show {**802.11a** | **802.11b**}

Information similar to the following appears:

```
802.11a Network.................................. Enabled
11nSupport....................................... Enabled
802.11a Low Band.......................... Enabled
802.11a Mid Band.......................... Enabled
802.11a High Band......................... Enabled
...
Pico-Cell-V2 Status.............................. Disabled
TI Threshold..................................... -50
Legacy Tx Beamforming setting................. Enabled
```

Step 7 See the beamforming status for a specific access point by entering this command:

show ap config {**802.11a** | **802.11b**} *Cisco_AP*

Information similar to the following appears:

```
Cisco AP Identifier.............................. 14
Cisco AP Name.................................... 1250-1
Country code..................................... US  - United States
Regulatory Domain allowed by Country............. 802.11bg:-A     802.11a:-A
...
Phy OFDM parameters
Configuration ............................... AUTOMATIC
Current Channel ............................. 149
Extension Channel ........................... NONE
Channel Width.............................. 20 Mhz
Allowed Channel List....................... 36,40,44,48,52,56,60,64,100,
..................................... 104,108,112,116,132,136,140,
..................................... 149,153,157,161,165
TI Threshold ................................ -50
```

```
Legacy Tx Beamforming Configuration ....... CUSTOMIZED
Legacy Tx Beamforming ..................... ENABLED
```

Commands Related to ClientLink

The following commands are related to ClientLink:

- The following commands are to be entered in the AP console:
  - To check the status of Beamforming on the AP, enter the **show controller d0/d1** command.
  - To find a client in the AP rbf table, enter the **show interface dot110** command.
  - To check the Beamforming rate assigned on the AP, enter the **debug d0 trace print rates** command.
- The following commands on the AP console are used for troubleshooting:
  - To show that ClientLink is enabled on a radio, enter the **show controllers | inc Beam** command.

    The output is displayed as follows:

    ```
    Legacy Beamforming: Configured Yes, Active Yes, RSSI Threshold -50 dBm
    Legacy Beamforming: Configured Yes, Active Yes, RSSI Threshold -60 dBm
    ```

  - To show that ClientLink is Beamforming to a particular client, enter the **show interface dot11radio 1 lbf rbf** command.

    The output is displayed as follows:

    ```
    RBF Table:
    Index      Client MAC      Reserved      Valid    Tx BF     Aging
     1       0040.96BA.45A0      Yes          Yes      Yes       No
    ```

Controller Planning

The following items affect the number of controllers required in a mesh network:

- Mesh access points (RAPs and MAPs) in the network.

  The wired network that connects the RAP and controllers can affect the total number of access points supported in the network. If this network allows the controllers to be equally available to all access points without any impact on WLAN performance, the access points can be evenly distributed across all controllers for maximum efficiency. If this is not the case, and controllers are grouped into various clusters or PoPs, the overall number of access points and coverage are reduced.

- Number of mesh access points (RAPs and MAPs) supported per controller. See Table 5: Mesh Access Point Support by Controller Model , on page 9.

  For clarity, nonmesh access points are referred to as *local* access points in this document.

Table 5: Mesh Access Point Support by Controller Model

Controller Model	Local AP Support (nonmesh)[1]	Maximum Possible Mesh AP Support
5508[2]	500	500
2504[3]	50	50
WiSM2	500	500

1 Local AP support is the total number of nonmesh APs supported on the controller model.
2 For 5508, controllers, the number of MAPs is equal to (local AP support - number of RAPs).
3 For 2504, controllers, the number of MAPs is equal to (local AP support - number of RAPs).



Note Mesh is fully supported on Cisco 5508 Controllers. The Base License (LIC-CT508-Base) is sufficient for indoor and outdoor APs (AP152X). The WPlus License (LIC-WPLUS-SW) is merged with the base license. The WPlus License is not required for indoor mesh APs.